STATEMENT OF PROF. KEVIN FU, PH.D.

DEPARTMENT OF
ELECTRICAL ENGINEERING & COMPUTER SCIENCE
UNIVERSITY OF MICHIGAN
ANN ARBOR, MI

KNOWLEDGE-BASED AUTHENTICATION (KBA)

SUBMITTED TO THE
U.S. SENATE COMMITTEE ON HOMELAND SECURITY
AND GOVERNMENTAL AFFAIRS

HEARING ON
THE IRS DATA BREACH:
STEPS TO PROTECT AMERICANS' PERSONAL
INFORMATION

TUESDAY, JUNE 2, 2015

1 Introduction

Good afternoon, Chairman Johnson, Ranking Member Carper, and distinguished members of the Committee. I am testifying before you today on the use of instant "secret questions" in knowledge-based authentication (KBA) related to the recent IRS breach. I will explain the key properties of instant KBA to give you a better understanding of current challenges and vulnerabilities. I will close with recommendations on what can be done in the future to avoid similar, large-scale breaches.

My name is Dr. Kevin Fu. I represent the cybersecurity research community. Cybersecurity researchers innovate technologies and principles to improve cybersecurity as well as break security systems to understand their weaknesses and limitations. I am Associate Professor of Computer Science & Engineering at the University of Michigan where I teach and carry out research on how to improve the trustworthiness of computer systems. My educational qualifications include a Ph.D., master's degree, and bachelor's degree from M.I.T.'s Department of Electrical Engineering and Computer Science. We teach programming to over 1,300 undergraduates each year, and we teach a rigorous course in computer security to 440 undergraduates each year.

I am speaking today as an individual. All opinions, findings, and conclusions are my own and do not necessarily reflect the views of any of my past or present sponsors or employers.

2 Authentication

Authentication is the process of verifying a claimed identity. For instance, the IRS web site attempted to verify the claimed identity of a user before disclosing transcripts of tax returns. There are three basic means to authenticate an identity:

- Inheritance-based: Something you are (e.g., biometrics, fingerprints, iris scans, signatures)
- Ownership-based: Something you have (e.g., ID cards, mobile phones, tokens)
- Knowledge-based: Something you know (e.g., **secret questions**, passwords, PINs)

Let me focus on "secret questions," which is a form of knowledge-based authentication (KBA).

3 Knowledge-Based Authentication (KBA) with Secret Questions

There are two popular ways of using secret questions to authenticate a user: static and instant. The recent IRS "get transcript" breach involved an attack on the instant KBA. To better understand instant KBA, let me first contrast it with static KBA.

Static KBA. Users will recognize that many financial web sites and cloud services ask for answers to "secret questions" during the initial creation of an account. The secret questions serve as a backup mechanism to reset lost or forgotten passwords. Static KBA would not be appropriate for establishing *initial* trust in an identity. IRS did not use static KBA, and that is an appropriate design choice for their situation. Common static KBA questions include:

- Where were you born?
- What is your favorite food?
- What is your favorite sports team?
- Where did you meet your spouse?

Instant KBA. Instant KBA (also known as dynamic KBA) also uses secret questions to authenticate a user. However, in instant KBA, the user does not file answers to secret questions beforehand. Instead, the web site presents personal questions with answers readily available or purchasable from credit reports and other financial sources. Thus, the site can verify that the user knows certain knowledge that is more difficult to obtain than publicly available data. The IRS "get transcript" service used instant KBA questions to authenticate tax payers downloading transcripts of their tax returns. The questions and answers were provided by a third party from the private sector.

To verify your identity, please select your previous address:

A. 52 Church St

B. 1600 Pennsylvania Ave

C. Gettysburg

D. None of the above

Figure 1: An illustrative example of a hypothetical instant KBA question drawn from financial records to establish faith that a user is who they claim to be.

4 Anatomy of Instant KBA at IRS Get Transcript and SSA mySSA

My understanding is that multiple federal sites make use of private sector services for instant KBA secret questions to verify the authenticity of a claimed tax payer identity. For instance, the mySSA site[1] advertises that it uses Experian to verify claimed identities with four secret questions. I believe that IRS used a similar service.

[1] https://secure.ssa.gov/

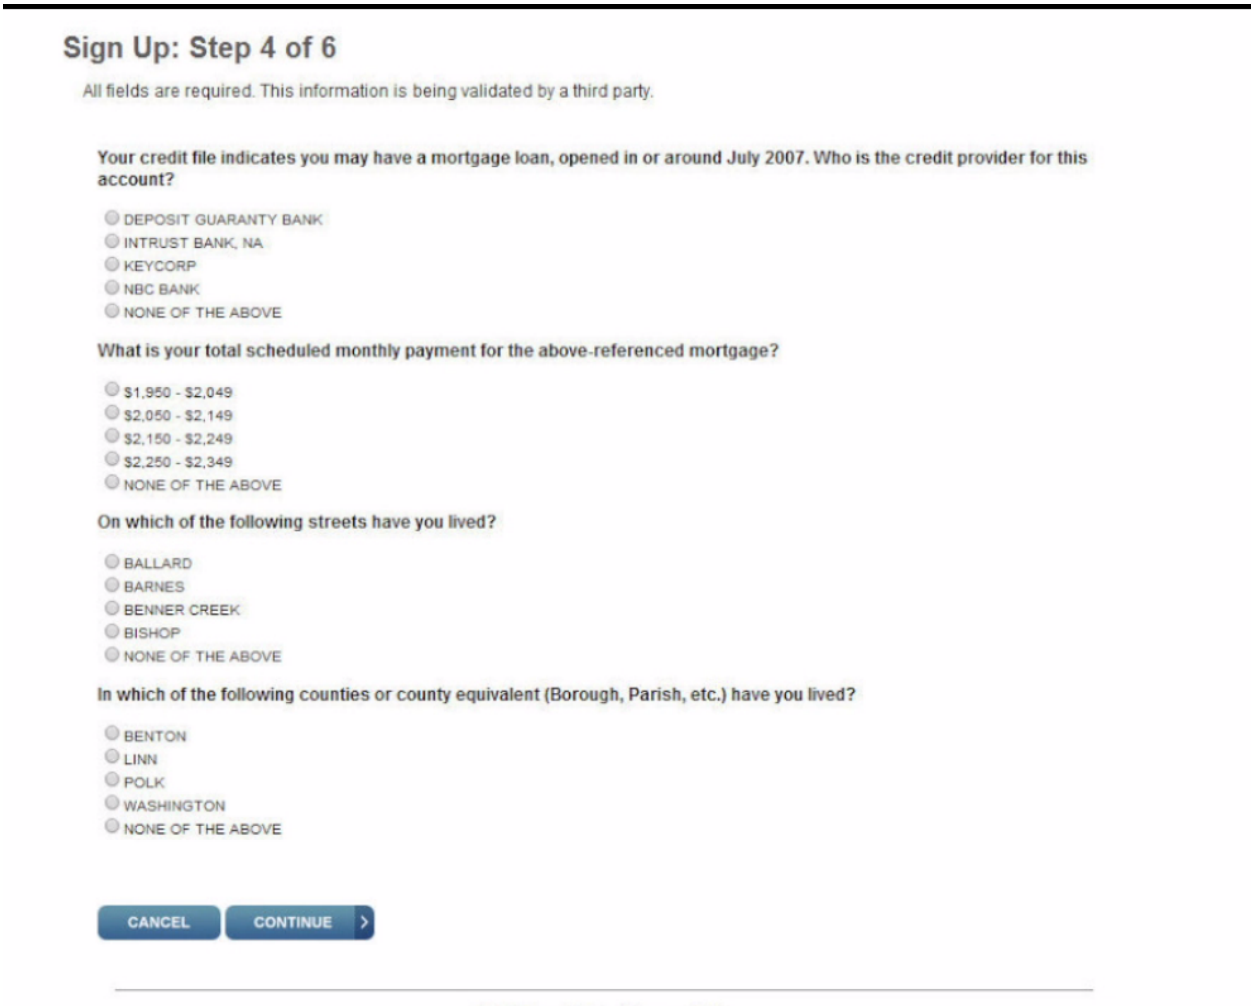

Figure 2: A screenshot of sample instant KBA secret questions at the IRS "get transcript" web site after having entered personal data and completed an email confirmation check from http://lpc.financialaidtv.com/cats/general_information#playlist-8075%3Avideo:video-5. It is believed that hackers answered these secret questions correctly by using personal data taken from any of the many breaches in the private sector.

5 KBA: Strengths and Limitations

No one system is perfect. Instant KBA does improve the security of identity verification by making it more difficult for an adversary to compromise a system, but sophisticated adversaries can nonetheless circumvent the protections at unprecedented scale, as demonstrated by the recent breach of 100,000 tax payer records at IRS. The root cause is that our supposedly independent systems are highly dependent on each other, and a seemingly unrelated compromise at one provider (e.g., Anthem, Target) can affect the security at a different service provider (IRS).

The main strength of instant KBA is ease of use. Most legitimate tax payers will be able to authenticate by answering multiple-choice questions about their personal, financial, and tax history. However, a major limitation is that the security of the system rests on assumption that the adversary does not have access to this information.

Instant KBA is designed under the threat model where an adversary may have stolen a tax payer's wallet. Using only the stolen wallet, it would be difficult for a criminal to answer four instant KBA questions successfully. Unfortunately, this threat model is no longer realistic as countless databases of such personal information have been compromised.

Opting out. Tax payers get no chance to opt out of the risks of instant KBA. As NIST explains in a technical report, "Instant KBA is not acceptable when transactions result in the release of sensitive or private information related to an individual." NIST researchers further explain:

> "In an Instant KBA authentication system, no matter how carefully the verifier treats the private personal authentication information, unless that information is known only to the verifier, the authentication system is somewhat at the mercy of third parties who may also have this information (and from whom the verifiers may have obtained it). But the user is initially not even a knowing party to the system, has given no consent, and has no obligation to treat every bit of personal private information that might be used in such a system as a secret, nor does the user know what personal private information may be used for authentication. It is inappropriate to involuntary expose the privacy of unknowing citizens to the risks of an instant KBA authentication scheme, unless the

risks for any individual citizen is very close to zero, however much an adversary may desire the information about that particular user. The (involuntary) users whose information is to be accessed, whether movie stars, public figures, or average citizens, may be expressly targeted by capable, experienced, resourceful attackers such as investigative reporters, private investigators, or personal enemies, who may be motivated to do a great deal of research to learn more about their target. Viewed from that perspective, instant KBA, with its vulnerability to off-line research, is more than a little alarming."

A principle espoused by the security research community is that if a user does not know something is a secret, then it's not a good secret for authentication because the user is easily tricked into divulging such information. Static and instant KBA therefore can violate this principle of security. In the 1990s, cybersecurity researchers initially believed that secret questions would be more secure than passwords. However, subsequent research in social and behavioral studies have shown severe weaknesses in authentication based solely on secret questions.

Lack of instant audibility. When I log into my Apple iTunes account from a new device, Apple sends me an email warning because Apple already knows how to reach me. When I make a credit card purchase, I receive a text message warning from my bank because my bank already knows how to reach me. With instant KBA, a service provider like IRS has no effective way to quickly inform a tax payer that their data or account was accessed because IRS does not already know how to reach the tax payer electronically. Worse, it is difficult for a tax payer to repudiate a fraudulent transaction.

Static KBA as a single factor is unreliable and insecure. Researchers at Google analyzed hundreds of millions of static KBA questions and have come to the conclusion that secret questions have poor reliability and poor security.[2] For this reason, Google services prefer to pair KBA ques-

[2] "Secrets, Lies, and Account Recovery: Lessons from the Use of Personal Knowledge Questions at Google" by Bonneau et al. in *Proceedings of the 24th International Conference on the World Wide Web*, May 2015. https://cdn.elie.net/publications/secrets-lies-and-account-recovery-lessons-from-the-use-of-personal-knowledge-questions-at-google.pdf

tions with second factors such as email or SMS messages to mobile phones. Note that the IRS *did* use email as a second factor; thus the adversary appears to have circumvented what Google recommends as a second factor. Their research paper lists a number of technical alternatives to personal knowledge questions in Section 6.2.

Static KBA questions are predictable. In 2009, researchers from Microsoft and Carnegie Mellon University (CMU) conducted a human subjects study of guessability of secret questions from static KBA. The study found that some secret questions had a 15% chance of guessability within five tries without knowing anything about the victim. In fact, the research foreshadowed the IRS breach with their warning that, "While most well publicized attacks on personal questions have been targeted at individuals, our results show that large scale attacks are also possible."[3] The authors recommend eliminating questions that are statistically guessable more than 10% of the time, and flagging questions that exceed a certain threshold of popularity.

6 Alternative Approaches

Let me highlight a few approaches that might improve the effectiveness of the authentication systems at IRS and other federal agencies.

Second-Factor Authentication (2FA). Use of a second factor can make it more difficult for an adversary to impersonate a tax payer online by slowing down or deterring attacks. A second factor should come from the "something you have" or "something you are" categories to be a genuine and independent second factor to "something you know." A popular second factor is a mobile phone. For example, DuoSecurity.com provides a suite of 2FA tools to combat credential theft and breaches. Federal service providers such as IRS could send a text message to a mobile phone to make it more difficult for a single adversary to impersonate 100,000 tax payers. However, no system is fool proof. It merely reduces risk. The threat landscape can change quickly.

[3] "It's No Secret. Measuring the Security and Reliability of Authentication via 'Secret' Questions" by Schechter et al. in *IEEE Symposium on Security and Privacy*, May 2009. http://research.microsoft.com/pubs/79594/oakland09.pdf

Notification warnings. One could imagine the IRS notifying a tax payer when someone attempts to access tax transcripts electronically. For instance, IRS could use contact information in tax returns to reach out to the tax payer or accountant to warn of the attempted download before allowing the download. But such systems are subject to phishing attacks, and would remove the instant gratification of quickly downloading tax returns.

NSTIC. The National Institute of Standards and Technology (NIST) launched the National Strategy for Trusted Identities in Cyberspace (NSTIC) for a ten-year goal of improving authentication of identities. The NSTIC roadmap set guiding principles for federal agencies to improve authentication by partnering with industry service providers. NIST published a report on its NSTIC interactions with IRS.[4]

In April 2015, NIST researchers involved with the NSTIC explained that, "While KBA is widely used today, there is no performance standard for KBA solutions—something that many of the NSTIC pilots have flagged as a significant challenge."[5]

Voice-Based Fraud Detection. The financial sector has been subject to widespread fraud by callers who attempt to engage in identity theft. One novel approach is to analyze the subtle cues in the audio conversation to identify known fraudsters. For instance, one might be asked to respond to instant KBA by phone rather than by typing. The subtle cadence and mannerisms of the speaker as well as the fundamental characteristics of the phone line makes it harder for an adversary to impersonate 100,000 people at once. The machine learning and security analytics are sufficiently effective that service providers can identify when a single fraudster calls back attempting to impersonate yet another consumer. This research was published in 2010, and later commercialized as a company called PinDrop.[6] There is also other research on speaker identification that may help with fraud detection.

[4] http://www.nist.gov/director/planning/upload/report13-2.pdf

[5] http://nstic.blogs.govdelivery.com/2015/04/09/a-retrospective-look-advancing-standards-for-strong-identity-and-authentication-in-the-identity-ecosystem/

[6] "PinDr0p: Using Single-Ended Audio Features To Determine Call Provenance" by Balasubramaniyan et al. in *ACM CCS*, 2010. http://www.cc.gatech.edu/ traynor/papers/traynor-ccs10.pdf and http://www.pindropsecurity.com/

7 Summary and Recommendations

The IRS used instant knowledge-based authentication in an attempt to verify identities seeking transcripts of tax returns. Unfortunately, the threat landscape is changing quickly as attackers adapt to newly fortified defenses. There will always be fraud, but a reasonable goal is to make it difficult for a single adversary to commit wide-scale, automated fraud. A major challenge in identity theft prevention is maintaining low false-positives (that would deny legitimate requests) and low false-negatives (that would allow identity theft) while serving the technologically diverse, tax paying U.S. population.

Technical recommendations include:

- Develop KBA performance standards and security metrics so that service providers such as IRS can more meaningfully decide acceptable risk of different kinds of KBA questions.
- Stop using SSNs or financial records as secrets for single-factor authentication because personal data are widely available in underground markets of stolen databases. Such data is effective only against unsophisticated adversaries.
- Consider enhancing KBA with a second factor of authentication from the “what you are” and “what you have” categories such as SMS messages or voice-based fraud detection.
- Leverage the existing cybersecurity expertise within the NIST’s National Cybersecurity Center of Excellence (CCoE), National Strategy for Trusted Identities in Cyberspace (NSTIC), and Information Security and Privacy Advisory Board (ISPAB).
- Encourage research collaboration between cybersecurity experts and social and behavioral science to carry out human subjects experiments that measure the risks and benefits of knowledge-based authentication.

We are likely to see more compromises of this nature because of systems depend on each other in subtle ways. What is most interesting is the advanced nature of the threat in the case of the IRS breach. Most cybersecurity research on the limits of secret questions does not consider the case when the adversary has a copy of the answers. Knowledge-based authentication systems are often built to protect against simple guessing attacks, but are not able to withstand an adversary with a complete cheat sheet of all the answers.

Let me end with an anecdote of an acquaintance affected by identity theft of their tax records in a previous incident. In April, an orthodontist received a notice via his CPA that someone had filed a fraudulent tax refund. To reclaim his identity, he had to fill out rather tedious affidavits. He now worries about the process he will have to follow for potentially the rest of his life to simply file a tax return. The identity theft protection does not make up for this orthodontist's significant time lost. Worse, his four-year-old child who was on one of the compromised tax returns may have to cope with the consequences of identity theft for the next hundred years of tax returns.

Thank you. I am happy to answer any questions you may have.

Biography

Kevin Fu is Associate Professor of Electrical Engineering and Computer Science at the University of Michigan. His cybersecurity research investigates how to achieve trustworthy computing on embedded devices with application to health care, commerce, and communication. He teaches computer science courses in security and privacy.

Prof. Fu received his Ph.D. in EECS from MIT where his research pertained to secure storage and how web authentication fails. Fu received a Sloan Research Fellowship, NSF CAREER award, Fed100 Award, and several best paper awards. Kevin was named MIT Technology Review TR35 Innovator of the Year. His participation in the provocative 2008 research paper analyzing the security of a pacemaker/defibrillator led to a wake-up call for cybersecurity in medical device manufacturing.

Kevin served as a visiting scientist on cybersecurity research at the U.S. Food & Drug Administration, the Beth Israel Deaconess Medical Center of Harvard Medical School, Microsoft Research, and MIT CSAIL. He is a member the NIST Information Security and Privacy Advisory Board. ISPAB is a Federal Advisory Committee that identifies emerging managerial, technical, administrative, and physical safeguard issues relative to information security and privacy in Federal Government information systems. Kevin is also the Chief Scientist and a co-founder of Virta Labs, Inc. whose product detects anomalies and malware with a smart power outlet rather than with installed software.

Please tell us about yourself

We collect and evaluate this information as a security measure to ensure that only you are able to access your personal information. We will not store your answers.

Why are these questions important?

You may have opened a mortgage loan in or around February 2013. Please select the lender to whom you currently make your mortgage payments. If you do not have a mortgage, select 'NONE OF THE ABOVE/DOES NOT APPLY'.

You may have opened an auto loan or auto lease in or around August 2013. Please select the lender for this account. If you do not have such an auto loan, select 'NONE OF THE ABOVE/DOES NOT APPLY'.

You may have opened a student loan in or around March 2004. Please select the lender that you have previously or you are currently making payments to. If you have not received student loans with any of these lenders now or in the past, please select 'NONE OF THE ABOVE/DOES NOT APPLY'.

According to our records, you graduated from which of the following High Schools?

. . .

Figure 3: A screenshot of sample instant KBA secret questions at the MySSA web site after having entered basic information such as a tax payer name, SSN, birthdate, and address. For each of the four questions, the user would select one of four multiple choices, often including an option of "None of the Above."

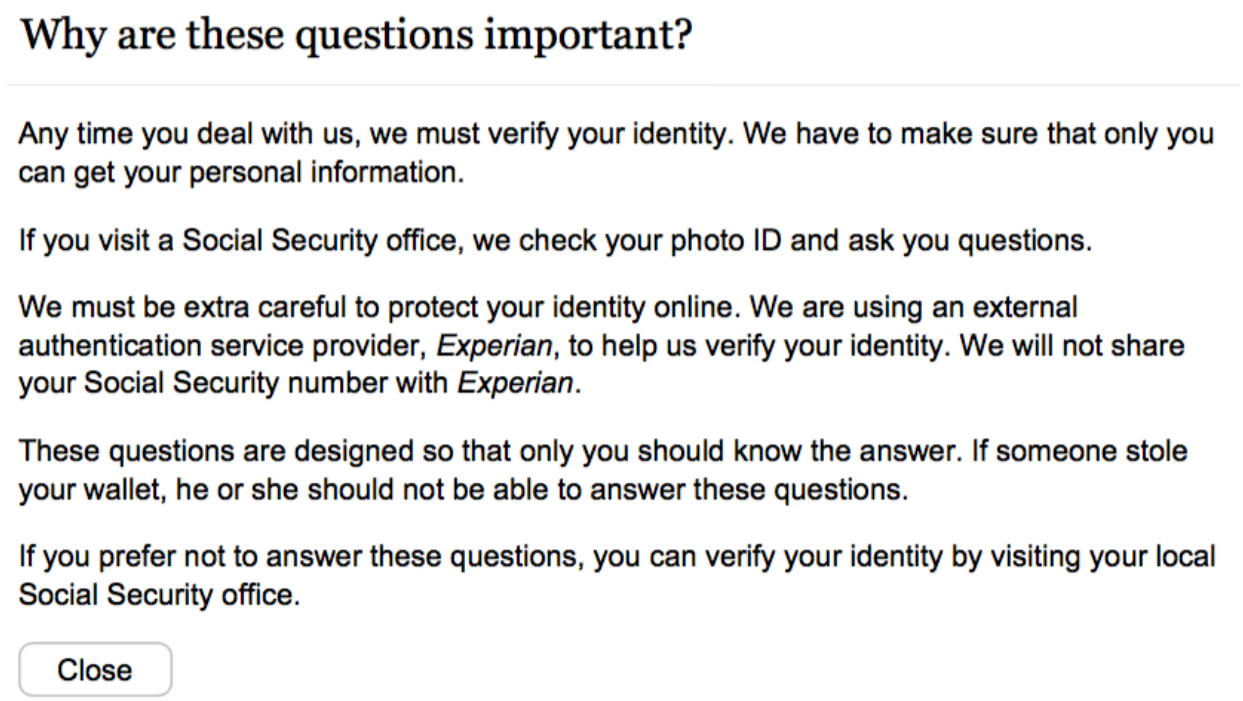

Figure 4: A screenshot of the MySSA web site that accurately explains their threat model. Unfortunately, the threats are changing quickly.